

ADA 112665

Semi-Annual Technical Report of Research

Under Contract No. F44620-73-C-0056,

Sponsored by Advanced Research Projects Agency

ARPA ORDER No. 2448

For the Period

1 November 1975 through 30 April 1976

CODING SYSTEMS AND THE COMPREHENSION OF INSTRUCTIONAL MATERIALS

ARPA Order: 2448

Program Code: 3D20

Name of Contractor: University of Oregon

Effective Date of Contract: 1 May 1973

Contract Expiration Date: 31 August 1976

Amount of Contract Dollars: \$332,417

Principal Investigator: Ray Hyman  
Professor of Psychology  
(503) 686-4910 or 686-4962

Program Manager: D. C. Howard  
Research Financial Administrator  
(503) 686-3146

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

APPROVED FOR PUBLIC RELEASE  
DISTRIBUTION UNLIMITED



DTIC  
ELECTE  
MAR 30 1982  
S D

D

82 03 30 074

DTIC FILE COPY

0.0. SUMMARY

The initial hardware problems with the Prime seemed to have been resolved when new problems were discovered within the central processor itself. The circuit board involved has been sent back to the manufacturer. We hope to have the Prime system ready to operate--both in terms of hard and software--in time for the summer.

Wickelgren and his students developed a new method--the "yes-no recall" method for studying retrieval dynamics of recall within the same paradigm that he has been using to study retrieval of recognition memory. One surprising finding is that although the ultimate strength of recognition memory is stronger than that of recall, the dynamics of retrieval--time to contact trace and rate of buildup of information--are the same for both. The speed-accuracy paradigm is being applied to the study of imagery, semantic memory, and speech recognition.

Reicher and Richards completed a series of studies on the representational requirements of high speed search. They were investigating the implications of an earlier finding that it is more difficult to search for a familiar target amongst unfamiliar background items than vice versa. The weight of their data strongly suggests that this is an attentional phenomenon--with unusual backgrounds the subjects have difficulty ignoring the irrelevant items. This finding is consistent with a finding that involves reading through prose passages. It may embody an important issue for skilled performance in any cognitive task. Successful performance involves the development of schemata that enables the performer to more quickly isolate and ignore irrelevant information.

Reicher has also developed, after many false starts, a way to experimentally control reading rate without destroying the "naturalness" of the reading task. He is planning a program of studies based on this procedure. The first experiments deal with the question of what is gained by a preliminary scan of the material prior to a more deliberate reading.

Hyman and Neill have found that their subjects can be quite ingenious in devising strategies to evade encoding stimulus materials according to organizations imposed by the experimenters. To get around this, Hyman has devised a content area for his subject matter (the location of cities on a map) that more naturally lends itself to a common form of representation for different subjects. Just how well this will succeed in inducing subjects to encode the material in predictable ways remains to be seen.

Hyman is also developing new lines of research to get at the basic problem of how prior memory organization aids and hinders the assimilation of new inputs. He has begun by collecting cases in which outstanding and intelligent men have been trapped into cognitive errors. The study of such cases has much been neglected. By studying only cognitive success stories we fail to learn what it is that distinguished the successes from the failures. Also, in agreement with Herb Simon, we believe that more of the cognitive apparatus shines through in times of failures than in times of successes.

P-1

## 1.0. INTRODUCTION

Because this is the last technical report prior to the completion of the project on August 31, 1976, we will report only some highlights of our work during the period under review. We will save a more detailed coverage for the final report which will be prepared upon completion of the contract.

### 1.1. Objectives

Our purpose on this contract was to close the gap between basic research on cognitive processes and the needs of instructional technology. The study of memory, perception, thinking, decision making and related processes should have obvious relevance to the design of instructional materials and procedures. But writers have consistently pointed out that experimental psychology has not been helpful for the design of instructional technologies as we might have hoped. One reason is a matter of scale in the time intervals and the amounts of materials employed in the laboratory as compared with those ordinarily dealt with in instructional settings. Courses of instruction range over weeks or years; laboratory experiments deal with material that is mastered within a single session or, at most, a few sessions, of approximately an hour or less. The materials employed in instructional courses are obviously meaningful and relevant to the future goals of the learner; the stimulus materials in laboratory experiments are typically meaningless and unstructured and have no obvious utility outside of the experiment itself. And the total amount of information in a typical course of instruction exceeds by many orders of magnitude the total amount of information in even the most demanding of laboratory experiments.

So one way to try to close the gap is to employ stimulus materials in the laboratory that more closely approximate in content, structure, and amount the sorts of information imparted in instructional courses. At the same time, the total number of sessions devoted to mastery of the materials can be increased so as to make the total time more equivalent.

Such changes in stimulus materials and total time spent on individual subjects, in turn, require new ways of specifying both independent and dependent variables as well as new experimental and analytical procedures to cope with the increased complexity of the design and analysis phases of the research.

Within this context our focus is upon questions relating the existing memory structures of the subjects to the type and format of the input materials. When we employ meaningful, as contrasted with meaningless, materials we increase the dependence of the outcome upon what the subject already knows or believes about the themes in the material prior to coming to the laboratory. This means that we have to devise procedures for assessing and taking into account the subjects' existing knowledge. We want to know how this existing knowledge and its organization constrains how the subject encodes and assimilates new information. At the same time we want to know how the new information--its

contents, organization, and relationships to what the subject already has in memory--affects or brings about changes in the pre-existing memory structure. Presumably, for instruction to be effective, some more or less permanent changes in the subject's memory structure must be achieved.

### 1.2. Research Plan

As the objectives imply, the research strategy is to perform experiments on cognition--encoding, storing, retrieving, and using information--employing both longer and more experimental sessions as well as stimulus materials that more closely approximate the content and structure of instructional materials. Much of our efforts involve trying out and discarding various types of materials; trying out and discarding various paradigms and experimental strategies; etc.

One example would be Reicher and his colleague's attempts to study reading within the laboratory situation. In earlier experiments some years ago, Reicher and Hyman attempted to see what it is, if anything, that speed readers were picking up from the text as they sped through it at enormous rates. One serious difficulty we encountered is that our readers could not be made to read at the fast speeds they had been trained to do in their speed-reading course. Without being conscious of the fact, they slowed down their speed dramatically, presumably because they knew they were going to be tested on the material. With this experience and others in mind, Reicher wanted to find a way to control the subject's reading pace without interfering with the "naturalness" of the process. Many, many ways were tried. Often the ones that appeared at first to be most natural turned out to be just the opposite. Eventually Reicher hit upon a method that appears to enable the subjects to read in a natural manner and yet gives the experimenter control over the pacing. Most of the experiments that were tried and failed in this series, of course, do not show up on the record as work accomplished.

## 2.0. WICKELGREN AND HIS ASSOCIATES

### 2.1. Recall vs. Recognition

Wickelgren and his coworkers developed a new method to facilitate comparison of memory retrieval in recall and recognition. In the "yes-no recall method", the subject, having previously learned an A-B association or both A-B and A-C associations, receives the A item and tries to generate (recall in his/her mind) the B or the B' and C items. After a variable lag (e.g., .2, .3, .5, .7, 1, 1.5, 2, or 3 sec), a second item is presented. This item may be a correct (B or C) item or an incorrect (D) item and the subject must make a "yes-no" matching decision within .4 second after the onset of the second item. From previous work, we know that the .4 second interval is too short for the subject to contact the memory trace of the second item (recognize it) from its physical presentation. Thus, this procedure enables Wickelgren to apply, for the first time, the speed-accuracy-tradeoff approach to the study of recall dynamics in addition to recognition dynamics.

A striking finding of the first experiment by Wickelgren and Corbett using this new procedure was that, while the asymptotic strength for recall was about one-half that for recognition, retrieval dynamics (rate and intercept parameters) were the same for both. Furthermore, (A-B, A-C) associative interference did not have a significantly greater effect on retrieval dynamics for recall than for recognition. According to Wickelgren, "The most parsimonious account of these data at present is that the same direct-access memory retrieval process operates in recall and recognition with parallel retrieval of all associations to a common stimulus item, but at a somewhat reduced rate the greater the level of competing associations".

## 2.2. Other Work by Wickelgren and Students

Within the framework and methodology described above, Wickelgren and his students are in the process of investigating aspects of trace-strength (direct access versus serial search), imagery (the role of mnemonics on the various parameters of memory dynamics), priming and short-term memory (transition of memory traces from passive to active and vice versa), semantic memory (concept-example associations using both natural and experimentally learned concepts), retrieval dynamics in propositional networks, speech recognition, and memory storage dynamics (consolidation and forgetting).

## 3.0. REICHER AND HIS ASSOCIATES

### 3.1. Representational requirements for high speed visual scanning

Reicher, Snyder, and Richards have submitted for publication a report on a series of nine experiments that demonstrate that it is much easier to look for an uncommon character embedded among common ones than to look for a common character among uncommon ones. Taken together, the experiments converge on the conclusion that it is the nature of the background items (the ones that the subject is attempting to ignore) rather than the nature of the target items that is more important to performance.

The initial impetus for these experiments arose out of some subjective and counterintuitive experiences by the investigators. But, as we will argue later on, the underlying phenomenon apparently has widespread generality across a variety of naturalistic tasks wherein individuals are trying to extract relevant information from an abundance of input. Reicher's findings seem to have implications, both theoretical and practical, for other research in this project--especially for some of the concerns of Hyman and his coworkers.

Some time ago, Reicher and Snyder were conducting a series of experiments to see if attention to a part of a visual display could be directed by "conceptual" (as opposed to simple physical) information in the stimulus. One task involved searching for a rotated letter in a matrix of upright letters. During some exploratory studies, these investigators completely inverted one of the stimulus matrices to see what looking for an upright letter among rotated ones would be like. They expected that it would be easier

than the other way around. They thought that the one letter in its familiar orientation would be very easy to find against a background of letters that were all in unfamiliar orientations. To their surprise, this task of finding the familiar among the unfamiliar turned out to be quite difficult--much more difficult than the reverse task of finding one unfamiliar form against a background of familiar ones. In the latter case, the subjective impression one gets is of the unfamiliar form standing out against the background--sometimes it literally seems to just pop out. In the reverse situation, the subjective impression is that the unfamiliar forms in the background keep competing for attention, making it difficult for the one familiar target to stand out.

Reicher, Snyder, and Richards describe their initial decision to study this phenomenon with the following justification:

"We believe that this is an interesting phenomenon in itself and also that it is theoretically interesting in several respects. It is evidence of directing attention by conceptual information... Further it seems to suggest one or both of two interesting possibilities. One is that people can ignore familiar but not unfamiliar items. The alternative is that unfamiliar items attract attention. Our subjective impressions were that the ability to ignore familiar but not unfamiliar items was the more important factor. What is particularly interesting about being able to ignore only well known characters is the implication that the act of ignoring this task requires well developed memory representations of the characters. The ability to make decisions about well learned information with little conscious involvement seems a requirement for a variety of perceptual skills. For example, highly skilled inspectors of manufactured goods may not consciously know why they rejected particular items unless there is some reason (and extra time) to look back... Chess masters show their ability to make better unconscious decisions in the observation that they consciously consider not more but better moves than lesser players... Reading might also involve unconscious monitoring of routine materials."

The basic experiments involved showing subjects a matrix with 9 characters. On target-present trials, eight of the characters were background items and one character was a target item. On target-absent trials, all 9 matrix positions were filled by background characters. The types of unusual characters employed were rotated English letters, partial letters and Gibson figures. The common characters were usually upright English letters but on one occasion they were digits. Search was easier through common backgrounds than through unusual backgrounds with all of the character types employed and whether measuring speed or accuracy.

Several alternative explanations might account for these findings. Some of the alternatives would make the findings uninteresting. For example, one possibility is that the subject merely examines each of the items in the matrix, one at a time in a sequential fashion. He identifies each item before moving on the next until he arrives at an item that matches the target he is seeking. The results, given this alternative, would be

explained on the basis that it takes longer to identify an unfamiliar item than it does a familiar one. The results become of more interest if it can be shown that we are dealing with an automatic, attentional phenomenon. The nine different experiments were designed to eliminate various alternatives and to pin down the most likely explanation.

Taken together, the set of experiments enables the authors to conclude the effect is not due to conscious sorting through of individual characters. From their findings, they draw the following implications:

"In many ways, people seem especially prepared to extract changes in stimulus information. For example, Sokolov's...work with habituation suggests that elaborate neural models of stimulus input can be generated, allowing the input to be monitored without attention as long as the stimulus does not change. Almost any change in the stimulus will cause an orienting response (or require attention). The neural model achieves its efficiency by performing routine memory checks without attention, presumably freeing the more executive functions for other higher-order processing. Thus, perceptual processes can be automated with an effect similar to that arising from automation of motor processes. The automation to which we refer here extends the range of perceptual automation into the realm of visual search and leaves open the possibility that a model similar to Sokolov's could be extended to handle a wide variety of perceptual skills requiring the rapid monitoring of routine input. Our present findings are of interest over and above the earlier finding that attention can be directed on the basis of conceptual categories. They imply that although subjects can ignore well learned characters, they cannot ignore characters because of an absence of substantial memory representations. The finding that automated accounting for stimuli requires well developed memory representations seems most concordant with the nature of skilled perceptual performance."

### 3.2. Further Implications of the Search Model

The effect found by Reicher and his colleagues is quite strong. They believe that it has widespread practical applications in the area of pattern recognition and reconnaissance. The general principle seems to be one that pervades all human perceptual and motor performance. Skilled performance depends upon developing highly practiced schema that can monitor both perceptual and motor activities at levels below conscious awareness. Such automation of relatively routine aspects of perception and performance frees the limited capacity of the conscious or central processor to cope with non-routine or unexpected intrusions into the ongoing routines. The big problem facing the novice or nonexpert in any field is knowing what he can ignore so that he can focus upon what is relevant. Part of this problem is having the necessary conceptual apparatus to segregate irrelevant from relevant.

As just one example of the generality of the principal that one can segregate out and ignore only that for which one has well-developed schemata, we can point to research done with quite different stimulus materials and within a completely different theoretical objective. Geiselman (Geiselman, R. E. Semantic positive forgetting: another cocktail party problem. Journal of verbal learning and verbal behavior., 1975, 14, 73-81) was interested in testing out hypotheses about cued forgetting with prose materials rather than unordered lists of words. Geiselman interweaved two prose passages, each with a different theme. His subjects read this intermixed material under instructions to remember only the material relevant to one of the themes and to forget the other material. Sometimes the sentences of the passage to-be-forgotten were scrambled, sometimes the sentences of the passage to-be-remembered were scrambled, and sometimes both passages were scrambled. When the target sentences were ordered rather than scrambled, this helped recall somewhat (an increment of approximately 1/10 of an additional sentence recalled out of a total of 10 sentences), but having the background sentences ordered rather than scrambled helped recall twice as much. As in the Reicher experiments, having a coherent background appears to be much more important in being able to ignore the background than having a coherent target. Among other interesting things, Geiselman's data also demonstrate a cost-benefit outcome--the gain in recall for items that subjects are supposed to remember is exactly cancelled out by the loss in recall of items they are supposed to forget.

### 3.3. Experiments on reading text

In all our work on this contract we face the very difficult challenge of finding tasks that simulate natural instructional material and demands. At the same time we need to be able to control and assess the strategies and procedures that our subjects employ in interacting with this material. Reicher and his associates have spent considerable time and effort in coming up with a good way to have people read text from a scope in a way which seems fairly natural to the subjects yet allows some control over the amount of time subjects are exposed to the material. In previous studies on reading and comprehension, Hyman and Reicher found that people will read much slower when they know they are to be tested. When we tried to employ different experimental conditions these became confounded with how fast the subjects attempted to read the material. So the problem is to control reading speed so that it is independent of other manipulations or independent variables that are to be studied.

Reicher and his associates tried a wide variety of display procedures. But all of them turned out to have undesirable qualities and distracting features. The best display that they have so far tried out seems promising. It allows most subjects to read in the range of natural reading rates. An entire paragraph is displayed on the scope and the subject is allowed to begin reading immediately. Two small arrows (one in each margin) pace the subject by showing him which row he should be on.

The first experiment employing this display procedure investigates what happens when a text is read twice. Many people have the impression that when



reading difficult material, a quick first reading (before a careful second reading) can be very helpful for understanding even though at the end of the first reading they can tell you almost nothing about the passage (this quick overview technique is actually a fixed component of many commercial speed reading and study programs). Whatever it is that this quick first scan leaves in memory, it does not seem to be anything that the reader can describe in any articulate manner. Perhaps the initial reading leaves the reader with a general idea of what the article is about (an anticipatory schema) and what its goals are so that facts can be better put into place on a subsequent reading. Although there is much talk of schemas aiding and sometimes foiling recall of textual material, it is not at all clear what a schema is or how it can be applied. A quite common notion of the usefulness of a schema is that it allows a person to get along without attending as much to subsequent presentations of a stimulus.\* The argument goes that as one learns more and more

---

\*In earlier work on speed reading and comprehension, Hyman and Reicher differed on what they thought a student of speed reading was learning to do. Hyman took the position that the student is not learning to literally read faster-- to somehow take in more information in a given interval of time. Instead, he was learning how to take advantage of redundancies. If he was reading an article in the Reader's Digest, say, he could take advantage of the fact that such articles usually consist of a statement of the theme at the beginning, several anecdotes to illustrate the theme, and then a statement of the conclusion at the end. Knowing this, he can speed through the article by simply reading the first paragraph and then skipping to the conclusion. He has not read faster; rather he has taken advantage of prior knowledge to read only what is necessary for a given purpose. Reicher, on the other hand, felt that some perceptual-motor skills could be learned such that the skilled reader actually could be taking in more information in a given unit of time. Almost surely, both of these viewpoints have some validity.

---

about an article or a field of knowledge he needs to read less and less subsequently because the memory representations that he has will fill in an incomplete picture. According to this view, the highly experienced individual would be less likely to notice or understand some unusual or unexpected statement (especially near the end of the article).

Another view of what a schema is borrows from Sokolov's views on the orienting reaction and habituation. This view suggests that with a lot of exposure to a topic, an individual can monitor with little attention much of the expected material thereby increasing his efficiency on the passage as a whole. An unexpected event would automatically draw the subject's attention. The schema in this case would enable the reader to better notice unusual or unexpected statements. Admittedly, someone with a lot of information about a subject matter would be able to make informed guesses about what might be written in the area of his expertise than an uninformed person. This means that the first view must be at least true in part.

The question here is whether there is anything more to schemas than this. More specifically, can a schema make possible more effective input processing. In the first experiments the subject reads through a passage (stories from the Reader's Digest) twice. He is tested with a recognition procedure that probes for ability to detect changes in text that do not and that do change the meaning. Reicher hypothesizes that subjects will better detect meaning changes as function of repetition but will not improve on ability to detect non-meaning changes (e.g., synonyms).

#### 4.0. HYMAN AND ASSOCIATES

##### 4.1. Selective Retrieval from an "implanted" data base

The work on selective versus automatic retrieval from a memory structure involves having the subject master a given data base and then testing for his ability to selectively retrieve information from that data base. Up to now our data bases have used as concepts a set of hypothetical individuals who were related to one another in two ways: 1) indirectly in terms of overlap in number of shared or common properties; and 2) directly, in terms of some explicitly stated relationship such as sibling, friend, parent, child, etc.

Our major problem or difficulty has been our inability to control strategies of representation that subjects employ. Often the subjects employ a mixed strategy. They will identify some of the hypothetical individuals with friends or acquaintances. Their subsequent reaction to such individuals in future tasks is then often determined by the properties of the acquaintance with which they have made the identification. The remaining hypothetical individuals are then organized according to one or more methods for dealing with objects and their lists of properties. Such mixed strategies are very messy and difficult for us to cope with experimentally. We, therefore, are trying to find other ways to arrive at internal representations that are "natural" and yet ones that we can deal with.

One current attempt is to employ data bases from real-world domains in which our intended organization is natural and likely to match the subject's way of organizing the material. We present subjects with a map of a hypothetical country upon which several towns are located. The subjects' task is to learn the location of the towns. When presented a blank map, for example, he has to be able to point to designated spots and name the town at that location. A variety of considerations suggest that we should succeed in creating memory organizations that are isomorphic to the spatial arrangement on the map. Next we could teach the subjects to store other properties for each town (size, whether on hill or in valley, etc.). Here we would predict that spatial location retrieval questions would behave according to the semantic distance findings in other research, but questions about common non-geographical properties will behave according to our current findings that questions about concepts close together in memory will be easier to answer regardless if they require an affirmative or negative response.

##### 4.2. Levels of Processing

According to Craik and Tulving, the levels of processing viewpoint has

created a new paradigm in memory research. It was introduced as an alternative to the information processing paradigm in which information was seen as flowing through several stages--sensory store, short-term store, long-term store, etc. The level's of processing approach presumably does away with fixed stages and tries to study memory in terms of what processing and to what depth the to-be-remembered material is exposed. As originally stated, however, this approach, although immediately fruitful, was full of ambiguities. From this writer's viewpoint, for example, the promulgators seem to confound such things as time devoted to processing, type of processing, number of different processes applied, and ultimate depth achieved. In its successful applications, adherents seem to find that when subjects are given tasks requiring attention to only the physical characteristics of words (capital or small letters), they remember the words worse than when given tasks requiring semantic processing (is it an animal or vegetable?).

Tram Neill cogently argues such findings are confounded with the fact that the memory tasks favor semantic rather than non-semantic retrieval. If we can devise plausible tasks in which the physical properties are needed for memory, then we very well reverse the findings that somehow semantic level is always deeper or superior to other levels. He has devised an experiment to start exploring this possibility.

#### 4.3. Other work by Hyman and coworkers.

The problem confronting our mission is to avoid getting bogged down by peculiarities of our specific laboratory tasks and the possibilities that they induce peculiar strategies that are not characteristic of real-life situations in which individuals commonly attempt to use information. When we first decided to see if the new trends in psycholinguistics and semantic memory could be applied to instructional materials, psychologists were beginning to argue that old memory research on nonsense syllables and unrelated lists of words was irrelevant because the basic unit in memory was the proposition--a concept and some property or relation specified for it. This led to much work on sentence memory instead of word memory. No sooner had this work gotten off the ground when it was attacked by Bransford and others who delighted in demonstrating that memory for single sentences was insufficient. Comprehension and memory often employed contextual and thematic material that was larger than sentences. These initial attacks were negative in that they did not specify what these larger-than-sentence units were or how to deal with them. Next, and the most recent phase, came the emphasis on frames, scripts, plots, and other higher order units within which subjects supposedly cope with meaningful materials. But even this new trend is being criticized for not going far enough to bring research and theory into consonance with natural cognition. Neisser, in a forthcoming book, argues that the frame concept does not go far enough; it is too static. He pushes for a concept of schema that is flexible and temporally extended as well as much larger than propositional units. The schema is what enables the individual to extract useful information from the environment. It is continually modified by the input, while at the same time directing the sensory and cognitive systems towards new inputs. It is anticipatory and flexible and generic in the sense that it can accomodate a wide variety of specific environmental inputs.

In our initial thinking on this project we believed that earlier work on memory was irrelevant because it dealt with meaningless and artificial stimulus materials and tasks. Subsequently, we have come to the conclusion that such material is quite relevant, but just how and in what way it is relevant cannot be decided without better information on how the individual copes with environmental information in extra-laboratory settings. Hyman has been pursuing this question in areas of cognitive success and failure. He has been studying technological and other innovations. He has also been studying outstanding cases of cognitive failures. He has been arguing that too much attention has been paid in general and technical education on providing us with models of successful achievements. But we also need to study cases in which outstanding thinkers go astray if we are to properly benefit from past experience and understand just how our cognitive structures operate.

One area that may be promising is to carefully study cases in which scholars and scientists have gone astray or wherein conflicts arise over the interpretation of information. The current interest in occult, supernatural and nonrational systems is one rich area for this study. For example, some physicist and engineers have become convinced through observations and experiments on the alleged psychic Uri Geller that this man can bend metal with his mind. At the same time, Hyman, through first hand and second hand contacts with Geller and these scientists is convinced that they have been badly deceived by a clever trickster. Both sides of the debate are adamant. Here we have just one case in which scientifically trained and accomplished individuals evaluate the same information and come to widely differing conclusions. Hyman has been examining cases in the past of similar disagreements among scholars. It turns out that such cases are much more widespread and common than is commonly believed. We ordinarily do not hear about them because they are embarrassing or because it is customary to talk about successes. But careful examination of such arguments teaches us much about how our cognitive systems operate and the situations under which they both succeed and can go astray. This in turn can help us use laboratory experiments to inform us of what we need to know to improve our ability to effectively employ our cognitive systems.